JMLR: Workshop and Conference Proceedings vol 40:1–23, 2015

MAX-CSP, Graph Cuts and Statistical Physics

Yi-Kai Wang WANGYIKAI1029@GMAIL.COM
State Key Laboratory for Novel Software Technology, Nanjing University

arXiv:1412.0340v3 [cs.DS] 14 Apr 2015

Abstract

The theoretical models providing mathematical abstractions for several significant optimization problems in machine learning, combinatorial optimization, computer vision and statistical physics have intrinsic similarities. We propose a unified framework to model these computation tasks where the structures of these optimization problems are encoded by functions attached on the vertices and edges of a graph. We show that computing MAX 2-CSP admits polynomial-time approximation scheme (PTAS) on planar graphs, graphs with bounded local treewidth, H-minor-free graphs, geometric graphs with bounded density and graphs embeddable with bounded number of crossings per edge. This implies computing MAX-CUT, MAX-DICUT and MAX k-CUT admits PTASs on all these classes of graphs. Our method also gives the first PTAS for computing the ground state of ferromagnetic Edwards-Anderson model without external magnetic field on d-dimensional lattice graphs. These results are widely applicable in vision, graphics and machine learning.

Keywords: Approximation Algorithm, Graphical Model, MAX-CSP, MAX-CUT, MAX-DICUT, MAX k-CUT, Computer Vision, Statistical Physics

1. Introduction

Given a graph $G(V, E)$, each vertex $i \in V$ represents a variable $\sigma_i \in [q]$ that can take q different values. Each vertex $i \in V$ is attached a function $\varphi_i : [q] \to \mathbb{R}$ and each edge $(i, j) \in E$ is attached a function $\varphi_{ij} : [q] \times [q] \to \mathbb{R}$. For a joint configuration of the vertices $\sigma \in [q]^V$, the energy of the configuration σ is defined as

$$E(\sigma) = \sum_{i \in V} \varphi_i(\sigma_i) + \sum_{(i,j) \in E} \varphi_{ij}(\sigma_i, \sigma_j). \tag{1}$$

Many significant problems in computer vision such as image segmentation Rother et al. (2004); Boykov and Funka-Lea (2006) and stereo vision Kolmogorov and Zabih (2002); Sun et al. (2003) are formalized as minimizing the energy functions defined on a lattice graph where each pixel represents a vertex in the lattice graph. The function $\psi_i(\cdot)$ measures the disagreement between the label and observed data at pixel i and $\psi_{i,j}(\cdot)$ measures the pairwise smoothness between pixel i and pixel j. The typical form of $\psi_i(\sigma_i)$ Boykov et al. (2001) is usually $\psi_i(\sigma_i) = (\sigma_i - p_i)^2$ where p_i is the pixel value of point i. The typical forms of $\psi_{i,j}(\sigma_i, \sigma_j)$ are Boykov et al. (1998)

$$\psi(\sigma_i, \sigma_j) = \begin{cases} w_{i,j} \cdot \big(1 - \delta(\sigma_i, \sigma_j)\big) & \text{(Generalized Potts model)} \\ \min\big(k_{i,j}, |\sigma_i - \sigma_j|\big) & \text{(Truncated absolute distance)} \\ \min\big(k_{i,j}, (\sigma_i - \sigma_j)^2\big) & \text{(Truncated quadratic distance)} \end{cases}$$

where $w_{i,j} \geq 0$ is a weight coefficient and $\delta(x) = 1$ when $x = 0$, and $\delta(x) = 1$ otherwise. $k_{i,j} > 0$ is a coefficient represents the maximum allowed penalty in the truncated model. When $q = 2$,

© 2015 Y.-K. Wang.

the energy minimization problem can be solved in polynomial time by maximum flow Greig et al. (1989) or graph cuts Boykov et al. (2001). However, in general this problem is still NP-hard. In this paper, we focus on the maximization of the energy function.

This model is equivalent to pairwise Markov random fields and factor graphs Kschischang et al. (2001) whose function nodes have exactly two adjacent variable nodes, which are extensively studied for probabilistic inference and maximum a posterior estimation in graphical models. Many significant optimization problems can be formalized by this model such as MAX 2-CSP and maximum graph-cut problems. A constraint satisfaction problem (CSP) is defined as a triple $\langle X, Q, C\rangle$ where $X = \{X_1, \ldots, X_n\}$ is a set of variables, $Q = \{Q_1, \ldots, Q_n\}$ is the set of respective domains of values of X, $C = \{C_1, \ldots, C_m\}$ is a set of constraints. The variable X_i can take the values in domain Q_i. Each constraint C_j is a function which takes a subset $S_j \subseteq X$ of variables as inputs and returns a boolean value as output representing whether it is satisfied. A CSP asks whether there is a configuration of variables such that all the constraints are satisfied. The most famous CSP is its boolean version SAT, whose objective is determining whether there is a truth assignment for the input CNF boolean formula.

MAX-SAT is the optimization version of SAT, which is significant for studying approximation algorithms and has been extensively investigated in the literature. Given a CNF boolean formula $\phi = \varphi_1 \wedge \varphi_2 \wedge \ldots \wedge \varphi_m$ over n boolean variables $X_1, \ldots, X_n$ where ϕ_j is a disjunction of literals for each $1 \leq j \leq m$, the objective of MAX-SAT is finding an truth assignment of the variables of ϕ such that the number of satisfied clauses is maximum. In its weighted version, each clause φ_j is assigned a weight w_j, the objective is finding an truth assignment of the variables of ϕ such that the sum of the weights of those satisfied clauses is maximum.

MAX-CSP generalizes CSP analogous to the way in which MAX-SAT generalizes SAT. In this paper, we focus on MAX 2-CSP where each constraint has exactly two input variables. For simplicity, we assume that each variable can take exactly q values, denoted by $[q] = \{0, \ldots, q-1\}$. The structure of each 2-CSP $\langle X, Q, C\rangle$ can be naturally represented by a graph $G(V, E)$ where each vertex $i \in V$ represents a variable $X_i \in X$. If X_i and X_j are the input variables of a constraint, then there is an edge $(i, j) \in E$ connecting vertex $i \in V$ and vertex $j \in V$. We set $\varphi_i = 0$ and $\varphi_{ij} : [q] \times [q] \to \mathbb{R}^+$ as

$$\varphi_{ij}(\sigma_i, \sigma_j) = \begin{cases} w_{ij}, & \text{if the corresponding constraint is satisfied} \\ 0, & \text{if the corresponding constraint is not satisfied} \end{cases}$$

where $\sigma_i \in [q]$ is the value of the variable represented by vertex i and w_{ij} is the weight of the constraint represented by edge $(i, j) \in E$. Then the objective of MAX 2-CSP is equivalent to computing $\max_{\sigma \in [q]^V} E(\sigma) = \max_{\sigma \in [q]^V} \sum_{(i,j) \in E} \varphi_{ij}$ where the configuration $\sigma \in [q]^V$ is a joint assignment of values to the variables represented by the vertices of G.

The graph-cut problems which are also widely applicable in learning algorithms have similar structures. Given an undirected graph $G(V, E, w)$ where $w : E \to \mathbb{R}^+$ assign a nonnegative weight w_{ij} to each edge $(i, j) \in E$. The goal of MAX-CUT problem is dividing the vertices of G into two sets S and $\bar{S}$ such that the value $\sum_{(i,j) \in C(S,\bar{S})} w_{ij}$ is maximum where $C(S, \bar{S})$ is the set of cut edges between S and $\bar{S}$. We set $\varphi_i = 0$ and φ_{ij} as

$$\varphi_{ij}(\sigma_i, \sigma_j) = \begin{cases} w_{ij}, & \text{if } \sigma_i \neq \sigma_j \\ 0, & \text{if } \sigma_i = \sigma_j \end{cases}$$

where $\sigma_i = 1$ if $i \in S$ and $\sigma_i = 0$ if $i \in \bar{S}$. Then computing the MAX-CUT is also equivalent to computing $\max_{\sigma\in[q]^V} E(\sigma) = \max_{\sigma\in[q]^V} \sum_{(i,j)\in E} \varphi_{ij}$. For MAX-DICUT problem whose input is a directed graph and whose goal is dividing the vertices of G into two sets S and $\bar{S}$ such that the total weight of the directed cut $\sum_{i\in S, j\in\bar{S}, (i,j)\in E} w_{ij}$ is maximum, the function φ_{ij} attached on the edge can be set as

$$\varphi_{ij}(\sigma_i, \sigma_j) = \begin{cases} w_{ij}, & \text{if } \sigma_i = 1 \wedge \sigma_j = 0 \\ 0, & \text{otherwise.} \end{cases}$$

We also exploit the MAX k-CUT problem, which is a natural generalization of MAX-CUT problem. Given a connected undirected graph $G(V, E, w)$ where $w : E \to \mathbb{R}^+$ assign a nonnegative weight w_{ij} to each edge $(i, j) \in E$, a k-cut is a set of edges $E' \subseteq E$ whose removal decomposes the input graph into k disjoint subgraphs. The goal of MAX k-CUT problem is computing such a set of edges E' that $\sum_{e\in E'} w_e$ is maximum. When $k = 2$, the MAX k-CUT problem is precisely the MAX-CUT problem. When $k > \Delta$ where Δ is the maximum degree of G, the optimal solution is precisely the sum of all the weights of the edges. Thus this problem is only interesting when $k \leq \Delta$.

This model also widely appears in statistical physics to provide mathematical abstraction for structure of materials. The most famous one is Potts model Potts (1952) which is used frequently in computer vision as we have mentioned. Another famous model is Edwards-Anderson model for describing the interactions of particles in spin glasses. The input graph is usually a d-dimensional lattice graph $\mathbb{L}^d$ where each vertex i in the lattice is a Ising spin $\sigma_i \in \{-1, +1\}$. The energy function Edwards and Anderson (1975) is $E(\sigma) = \sum_{(i,j)\in E} J_{ij}\sigma_i\sigma_j + B\sum_{i\in V} \sigma_i$ where the parameters J_{ij} are exchange couplings and $B\sum_{i\in V} \sigma_i$ is the external magnetic field. The interactions between spins σ_i and σ_j is ferromagnetic if $J_{ij} > 0$ and antiferromagnetic if $J_{ij} < 0$. The Gibbs measure of a given configuration $\sigma \in [q]^V$ is defined as $P(\sigma) = \frac{1}{Z(\beta)} \exp(-\beta E(\sigma))$ where the normalizing constant Z is the *partition function* that $Z(\beta) = \sum_{\sigma\in[q]^V} \exp(-\beta E(\sigma))$. For the ferromagnetic Edwards-Anderson model without external magnetic field that $J_{ij} > 0$ and $B = 0$, the energy function of the Edwards-Anderson model is $E(\sigma) = \sum_{(i,j)\in E} J_{ij}\sigma_i\sigma_j = C - 2\sum_{(i,j)\in E, \sigma_i\neq\sigma_j} J_{ij}$. where $C = \sum_{(i,j)\in E} J_{ij}$. In this case, the minimization of the energy function is equivalent to computing the MAX-CUT of the input graph, which is polynomial-time solvable on planar graphs. This problem is NP-hard even on 3-dimensional lattice graphs. Other optimization techniques such as branch and bound methods, belief propagation have also been applied to this problem but no provable bound has been achieved. In this paper, we prove the following theorem.

Theorem 1 *Given a graph $G(V, E)$ and a set of functions attached on the vertices and edges, if it satisfies $\varphi_i \geq 0$ and $\varphi_{ij} \geq 0$ for all $i \in V$ and $(i, j) \in E$, then there is a PTAS for maximizing the energy function on planar graphs, graphs with bounded local treewidth, H-minor-free graphs, geometric graphs with bounded density and graphs embeddable with bounded number of crossings per edge.*

Theorem 1 is based on a more general theorem (Theorem 5) which will be specified later. If all the constraints of the 2-CSP are polynomial-time verifiable, we naturally have the the PTASs for computing MAX 2-CSP, MAX-CUT and MAX-DICUT.

Corollary 2 *Computing MAX 2-CSP admits PTASs on planar graphs, graphs with bounded local treewidth, H-minor-free graphs, geometric graphs with bounded density and graphs embeddable with bounded number of crossings per edge.*

Corollary 3 *There are PTASs for computing MAX-CUT, MAX-DICUT and MAX k-CUT on planar graphs, graphs with bounded local treewidth, H-minor-free graphs, geometric graphs with bounded density and graphs embeddable with bounded number of crossings per edge.*

Proof Choosing a set of terminals $S = \{s_1, \ldots, s_k\} \subseteq V$, the configuration of s_t where $1 \leq t \leq k$ is fixed to $t \in [q]$. The vertices in graph $G' = G \backslash S$ are free variables. The functions attached on $(i, j) \in E$ are the same as those of MAX-CUT. Then computing the MAX k-CUT for fixed S is equivalent to computing $\max_{\sigma \in [k]^{V(G')}} \sum_{(i,j) \in E} \varphi_{ij}$. By Theorem 1, we have a PTAS to compute MAX k-CUT for fixed S. There are at most $P(n, k) = \frac{n!}{(n-k)!} \leq n^k$ possibilities for S, which is polynomial if k is fixed. By enumerating all these cases, we have a PTAS for computing the MAX k-CUT. ■

Corollary 3 directly gives a PTAS for computing the ground state of ferromagnetic Edwards-Anderson model without external magnetic field on these classes of graphs. For minimization of the energy function, we do not have such a general result as Theorem 1. Furthermore, we prove the following result.

Theorem 4 *There does not always exist PTAS for minimizing the energy function even on planar graph unless P = NP.*

This result implies that there does not always exist PTAS for computing the min-sum of unbalanced f_i on graphs with bounded local treewidth, H-minor-free graphs, geometric graphs and graphs with bounded crossing number per edge and more general classes of graphs unless P = NP. However, we show that if the functions φ_{ij} can be well partitioned so that the functions attached on the vertices satisfy the *balance* property, minimization of the energy function also admits PTAS on these classes of graphs. This gives PTASs for minimizations of many typical energy functions in computer vision. The related works are introduced in Appendix A. The definitions of basic concepts and introduction of our technical tools are included in Appendix B.

2. Approximation Schemes

We build a unified framework to formalize the optimization problems. Consider a graph $G(V, E)$, directed or undirected. For directed graphs, the neighbors of vertex $i \in V$ is defined as $N^-(i) = \{j : (j, i) \in E\}$ and $N^+(i) = \{j : (i, j) \in E\}$. For undirected graphs, $N^-(i) = N^+(i) = N(i)$. If G is undirected, then for each $i \in V$, $f_i : [q]^{N(i)} \times [q] \to \mathbb{R}$ take the configurations of $N(i) \cup i$ as inputs and return a real number as output. If G is directed, then for each $i \in V$, then $f_i : [q]^{N^-(i)} \times [q] \to \mathbb{R}$ take the configurations of $N^-(i) \cup i$ as inputs and return a real number as output.

Although our model is well-defined on directed graphs, we only need to consider undirected graphs. For a directed graph $G(V, E)$, its corresponding undirected graph is constructed by replacing all the directed edges by undirected edges and then replacing multiple undirected edges between pairs of vertices if exists to a single undirected edge. The function f_i is replaced by f_i' which takes configurations of $N^-(i) \cup N^+(i)$ as inputs but only the $N^-(i)$ part can affect the output of f_i' which is equal to the output of f_i.

For each function φ_{ij} attached on edge $(i, j) \in E$, let α_{ij} and α_{ji} be two constants satisfying $\alpha_{ij} \geq 0$, $\alpha_{ji} \geq 0$ and $\alpha_{ij} + \alpha_{ji} = 1$. We set $f_i = \varphi_i + \sum_{j \in N(i)} \alpha_{ij} \cdot \varphi_{ij}$, then computing the energy

function $E(\sigma)$ is equivalent to computing $\sum_{i\in V} f_i$. Such a set of constants $\{\alpha_{ij} : (i,j) \in E\}$ is called a partition of the functions attached on the edges. By this reduction, we only need to consider the functions f_i attached on the vertices.

Throughout this paper, we ignore the computation time of f_i when analyzing the computational complexity. We assume that they are all $O(1)$-time computable.

Theorem 5 *If there is a partition $\{\alpha_{ij} : (i,j) \in E\}$ of the functions attached on the edges such that $f_i \geq 0$ for all $i \in V$ where $|V| = n$, then there is a PTAS for computing $\max_{\sigma\in[q]^V} \sum_{i\in V} f_i$ on planar graphs, graphs with bounded local treewidth, H-minor-free graphs, geometric graphs with bounded density and graphs embeddable with bounded number of crossings per edge.*

It is easy to see that Theorem 1 is a corollary of Theorem 5. The proof of Theorem 5 is based on the following lemma.

Lemma 6 *Given a graph $G(V,E)$ with treewidth bounded by k, for any $U \subseteq V$ and any set of $f_i \geq 0$ defined on G, $\max_{\sigma\in[q]^V} \sum_{i\in U} f_i$ and $\min_{\sigma\in[q]^V} \sum_{i\in U} f_i$ can be computed in $O(q^{O(k)}n)$ time.*

Proof Since the input graph G has treewidth bounded by k, we use the algorithm of Bodlaender (1993) to construct a tree decomposition $T = (I, F)$ rooted at $r \in I$ with treewidth k for G in linear time (but exponential in k). For each $i \in I$, the subtree of T rooted at i is denoted by T_i. The set of vertices in T_i is denoted by V_{T_i}. The configurations of bag X_i where $i \in I$ is denoted by σ_{X_i}. Suppose the child nodes of $i \in I$ are $i_1, \ldots, i_d \in I$ and the parent node of $i \in I$ is $p_i \in I$.

Our algorithm is a dynamic programming from the leaves to the root. We enumerate the all the possible configurations of $X_i - X_{p_i}$ for each bag X_i. For root r, $X_r - X_{p_r} = X_r$. By the definition of tree decomposition, $X_{i_x} \cap X_{i_y} \subseteq X_i$ for $1 \leq x \neq y \leq d$. Therefore, $X_{i_t} - X_i$ for $1 \leq t \leq d$ are pairwise disjoint. Let $S^U_{i\setminus p_i}(\sigma_{i\setminus p_i})$ denote the max-sum of the f_i attached on vertices in $(U \cap V_{T_i}) - (X_{p_i} \cup \partial X_{p_i})$ with the configurations of vertices in $X_i - X_{p_i}$ being fixed to $\sigma_{i\setminus p_i}$. The set ∂X_{p_i} denotes the vertices adjacent to vertices in X_{p_i} but not in X_{p_i}. Note that $S^U_{i\setminus p_i}(\sigma_{i\setminus p_i})$ does not include the sum of f_i attached on vertices in ∂X_{p_i}. This is because their values are not fixed when the configurations of X_{p_i} are not given. The value of $S^U_{i\setminus p_i}(\sigma_{i\setminus p_i})$ can be computed by the following recurrence:

$$S^U_{i\setminus p_i}(\sigma_{i\setminus p_i}) = \max_{\{\sigma_{i_t\setminus i}\}} \left\{ \Gamma^{\sigma_{i\setminus p_i}}_{X_i - X_{p_i}} + \sum_{t=1}^{d} S^U_{i_t\setminus i}(\sigma_{i_t\setminus i}) \right\}$$

where $\Gamma^{\sigma_{i\setminus p_i}}_{X_i - X_{p_i}}$ is the sum of f_i attached on vertices in $(U \cap (X_i \cup \partial X_i)) - (X_{p_i} \cup \partial X_{p_i})$ when the configuration of $X_i - X_{p_i}$ is fixed to $\sigma_{i\setminus p_i}$. Then $\max_{\sigma\in[q]^V} \sum_{i\in U} f_i = \max_{\sigma_{r\setminus p_r}} S^U_{r\setminus p_r}(\sigma_{r\setminus p_r})$.

Note that this recursion only holds for the set of f_i arising from energy functions taking the form of Equation (1) since φ_{ij} is pairwise so that we can enumerate all the possibilities of each $\sigma_{i_t\setminus i}$ for each $1 \leq t \leq d$ respectively. Even for f_i arising from MAX 3-CSP, this recursion does not hold any longer. The min-sum can be computed in the same way.

For each bag X_i, since $|X_i| \leq k+1$ for all $i \in I$, $\sigma_{i\setminus p_i}$ has at most q^{k+1} possible values. We compute $S^U_{i\setminus p_i}(\sigma_{i\setminus p_i})$ for each $\sigma_{i\setminus p_i}$ at most once. To compute $S^U_{i\setminus p_i}(\sigma_{i\setminus p_i})$, we also need to know $X_i \cap X_{p_i}$ for each X_i. These can be computed in $O(kn)$ time if the vertices in each X_i are stored

in order or using data structures such as hash tables. Furthermore, we need to decide whether a vertex is in the set $X_i \cap X_{p_i}$ or $X_i \cap \partial X_{p_i}$ when computing $S^U_{i\setminus p_i}(\sigma_{i\setminus p_i})$. The time complexity of the deciding whether a vertex is in $X_i \cap X_{p_i}$ is $O(1)$ if using hash tables. Whether a vertex $i \in V$ is in $X_i \cap \partial X_{p_i}$ can be decided by checking whether all the inputs of f_i are prepared. Since we have assumed that the computation time of f_i is $O(1)$, the time complexity is also $O(1)$. Therefore, the total time complexity of our dynamic programming algorithm is $O(q^{O(k)}n)$. ■

This algorithm is denoted by $\mathcal{DPT}_{G,\mathcal{F}}(U)$ where $\mathcal{F}$ is the set of f_i attached on G. In the rest of this section, we prove Theorem 5 for each class of graphs.

2.1. Planar graphs

Theorem 7 *If $f_i \geq 0$ for all $i \in V$ where $|V| = n$, then for fixed k, there is an $O(q^{O(k)}kn)$-time algorithm for computing $\max_{\sigma\in[q]^V} \sum_{i\in V} f_i$ that achieves a solution of value at least $k/(k+2)$ optimal for general planar graphs.*

Proof Given a planar embedding of a planar graph G, we decompose it into several disjoint $(k+2)$-outerplanar subgraphs by deleting all the edges between levels congruent to $\ell \bmod (k+2)$ and $\ell+1 \bmod (k+2)$ for some integer ℓ that $0 \leq \ell \leq k+1$. A possible partition of an 8-outerplanar graph with $k=3$ and $\ell=4$ is shown below in Figure 1. The deleted edges are represented by red dash lines and the boundary nodes (level 4 and level 5) and boundary lines are colored by orange. The disjoint $(k+2)$-outerplanar subgraphs are denoted by $G_1, G_2, \ldots, G_t$ from outside to inside. The corresponding vertex sets are denoted by $V_1, V_2, \ldots, V_t$. Each G_i where $1 < i < t$ consists of exactly $k+2$ levels of nodes and is surrounded by vertices of the inside boundary and the outside boundary. The number of levels of G_1 and G_t may be less than $k+2$. G_1 may have only inside boundary vertices and G_t may have only outside boundary vertices. To unify the notations, each G_i is a k_i-outerplanar graph where $1 \leq k_i \leq k+2$.

The set of outside boundary nodes of G_i is denoted by $V_{i,1}$ and the set of inside boundary nodes of G_i is denoted by V_{i,k_i}. The set of remaining nodes is denoted by A_i where $A_i = \bigcup_{j=2}^{k_i-1} V_{i,j}$. The set of boundary vertices of G_i is denoted by $B_i = V_i - A_i$. We further let $A = \bigcup_{i=1}^{t} A_i$ and $B = \bigcup_{i=1}^{t} B_i$. Each G_i is a k_i-outerplanar graph where $k_i \leq k+2$, thus the treewidth is bounded by $O(k)$. We construct an optimal tree decomposition T_i for each G_i in linear time using the algorithm of Bodlaender (1993). Let $\mathcal{F}_i$ denote the set of f_i attached on vertices of G_i. For each G_i, we use the dynamic programming algorithm $\mathcal{DPT}_{G_i,\mathcal{F}_i}(A_i)$ to maximize the sum of f_i attached on the vertices of A_i while ignoring the values of f_i attached on B_i. By Lemma 6, the dynamic programming on G_i for all $1 \leq i \leq t$ can be computed in parallel with time complexity $O(q^{O(k)}n)$.

Suppose the optimal solution is σ_{OPT} and the max-sum of f_i is S_{OPT}. The sum of f_i attached on A_i computed by the dynamic programming is denoted by S_{A_i}. The sum of f_i of $A_1, \ldots, A_t$ computed by dynamic programming is S_A. The sum of f_i on A_i is denoted by S_i when the configurations of A_i are fixed to σ_{OPT}. We have $S_{A_i} \geq S_i$. Therefore it holds that $S_A = \sum_{1\leq i\leq t} S_{A_i} \geq \sum_{1\leq i\leq t} S_i$. By pigeon hole principle, for at least one ℓ, the sum of the f_i attached on boundary nodes B is at most $\frac{2}{k+2}$ of S_{OPT}. Since $f_i \geq 0$ for all $i \in V$, whatever the configurations of the boundary nodes are, the sum of the f_i on boundary nodes is non-negative, denoted by S_B. Therefore it holds that $S_A + S_B \geq \sum_{1\leq i\leq t} S_i \geq \frac{k}{k+2} \cdot S_{\text{OPT}}$. Therefore our algorithm gives a solution that is at least $\frac{k}{k+2}$ optimal. For each $0 \leq \ell \leq k+1$, we need to repeat the dynamic programming of $O(q^{O(k)}n)$ time.

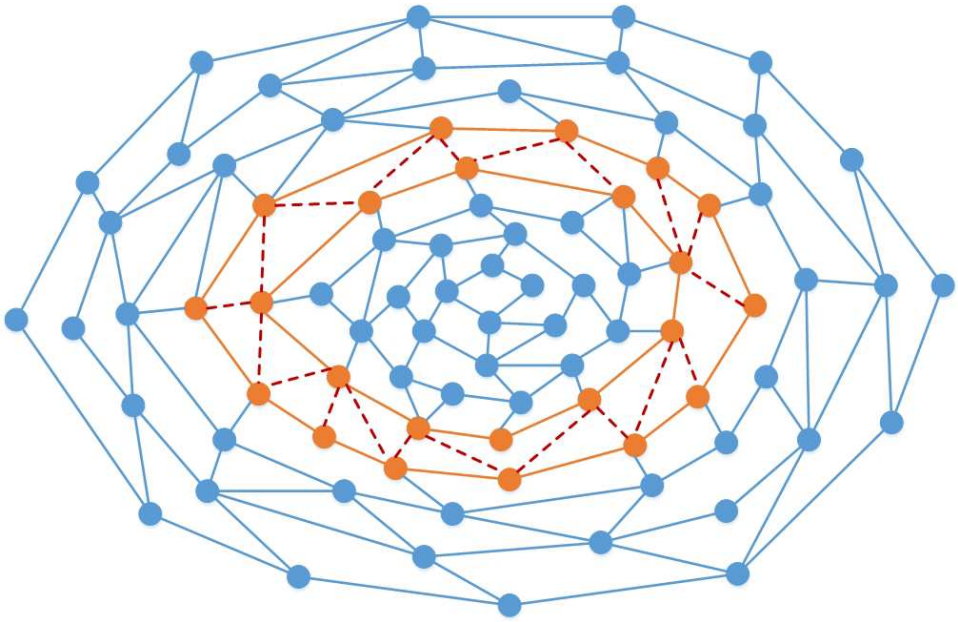

Figure 1: An 8-outerplanar graph with a possible partition

The total time complexity is $O(q^{O(k)}kn)$. This completes the proof. ■

It implies that if $f_i \geq 1$ for all $i \in V$, then the value of $\max_{\sigma\in[q]^V} \prod_{i\in V} f_i$ can be approximate to $(\max_{\sigma\in[q]^V} \prod_{i\in V} f_i)^{k/(k+2)}$ in $O(q^{O(k)}kn)$ time on planar graphs. This algorithm for approximating the max-product of f_i is not a PTAS. Its performance depends on the value of the maximum product and k. The approximation ratio of our algorithm when $k = 18$ is shown in Figure 2. Similar bounds also hold on graphs with bounded local treewidth, H-minor-free graphs, geometric graphs with bounded density and graphs with bounded crossings per edge which will be specified later. Our technique can serve as an alternative of classic methods such as max-product or min-sum belief propagation for MAP estimation. Furthermore, to our knowledge, such provable bounds have not been known for belief propagation Pearl (1988) and its generalized versions Yedidia et al. (2000).

Definition 8 *The functions $f_i \geq 0$ defined on graph $G(V, E)$ are α-balanced if for each $i \in V$ there is a constant b_i that $b_i \leq f_i \leq \alpha \cdot b_i$ where $\alpha \geq 1$ is a constant. The number $b_i \geq 0$ is called the balancer of vertex $i \in V$.*

If the functions f_i are balanced, we trivially achieve a $\frac{1}{\alpha}$-approximation algorithm for maximization and an α-approximation algorithm for minimization of the energy function. When the functions f_i are balanced, we have the following result for computing the min-sum.

Theorem 9 *If $f_i \geq 0$ is α-balanced, then for fixed k, there is an $O(q^{O(k)}kn)$-time algorithm for computing $\min_{\sigma\in[q]^V} \sum_{i\in V} f_i$ that achieves a solution of value at most $1 + \frac{2\cdot(\alpha-1)}{k+2}$ optimal for general planar graphs.*

The proof of Theorem 9 is similar to the proof of Theorem 7. The only difference is we use $\mathcal{DPT}_{G_i,\mathcal{F}_i}(A_i)$ to minimize the sum of f_i on A_i.

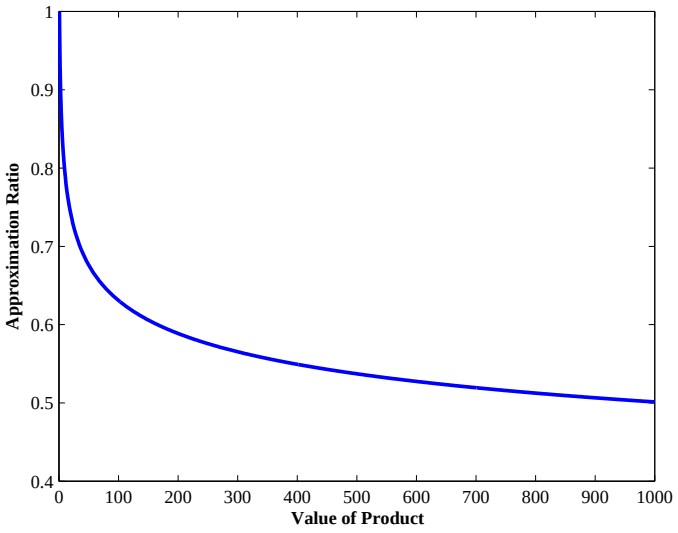

Figure 2: The approximation ratio when $k = 18$

Theorem 10 *There does not always exist PTAS for computing the min-sum of unbalanced f_i even on planar graphs unless P = NP.*

Proof We simply construct a reduction from computing the chromatic number to computing the min-sum of f_i. We set f_i as

$$f_i = \begin{cases} w_x & \text{if } \sigma_i = x \in [q] \text{ and } \sigma_i \neq \sigma_j \text{ for } \forall (i,j) \in E \\ +\infty & \text{if } \sigma_i = \sigma_j \text{ for any } (i,j) \in E \end{cases}$$

where w_x satisfies $w_x > n \cdot w_{x-1}$. This is allowed because both G and the set of functions f_i serve as inputs. By this setting, the functions f_i are not balanced. If the graph G is x-colorable, then the value of the min-sum $S_{\min} = \min_\sigma \sum_{i \in V} f_i$ must fall into the interval $[w_x, w_x \cdot n]$. For any $x, y \in [q]$, $[w_x, w_x \cdot n]$ and $[w_y, w_y \cdot n]$ are pairwise disjoint. By four color theorem, any planar graph is 4-colorable. Furthermore, it is known that 3-coloring problem remains NP-complete even on planar graph of degree 4 Dailey (1980). It implies that it is NP-hard to approximate the chromatic number within approximation ratio $4/3$ even on planar graphs. Therefore if we have a PTAS for computing the min-sum of f_i on planar graphs, then we have a polynomial time algorithm for computing the chromatic number of planar graphs, which leads to a contradiction. ■

By enlarging the gaps between the disjoint intervals, we can achieve stronger inapproximability results by this reduction.

2.2. Graphs with bounded local treewidth

Eppstein (2000) proves that a graph has bounded local treewidth if and only if it is apex-minor-free. Demaine and Hajiaghayi (2004b) prove that any apex-minor-free graph has linear local treewidth.

Theorem 11 *If $f_i \geq 0$ for all $i \in V$ where $|V| = n$, then for fixed k, there is an $O(q^{O(k)}kn)$-time algorithm for computing $\max_{\sigma\in[q]^V} \sum_{i\in V} f_i$ that achieves a solution of value at least $k/(k+2)$ optimal for graphs with bounded local treewidth.*

The proof of Theorem 11 is similar to the proof of Theorem 7. Choosing a vertex $v \in V$, we construct a BFS tree rooted at v. The level of a vertex is defined as its distance from v. We remove the edges between vertices whose levels are congruent to $\ell \pmod{k+2}$ and $\ell+1 \pmod{k+2}$. Then the input graph G is decomposed into several disjoint subgraphs with treewidth bounded by $O(k)$. We construct a tree decomposition for each subgraph and use the dynamic programming algorithm of Lemma 6 to maximize the sum of f_i attached on non-boundary vertices while ignoring the values of f_i on the boundary. By a similar argument, we can achieve the PTAS. For computing the min-sum, we have a similar result as Theorem 9. More details of the proof are specified in Appendix C.

2.3. H-minor-free graphs

A graph H is a k-apex graph of a graph G if $G = H\backslash A$ for some subset A of at most k vertices which is called *apices*. The definition of almost-embeddable graph is given as follows.

Definition 12 *A graph G is almost-embeddable on a surface Σ if G can be written as the union of $k+1$ graphs $G_0 \cup G_1 \cup \ldots \cup G_k$, satisfying the following conditions: (1). G_0 has an embedding on Σ. (2). The graphs $G_1, G_2, \ldots, G_k$ are pairwise disjoint, called vortices. (3). For each index $i \geq 1$, there is a disk D_i inside some face F_i of G_0, such that $U_i = V(G_0) \cap V(G_i) = V(G_0) \cap D_i$. Moreover, the disks D_i are pairwise disjoint. (4). For each index $i \geq 1$, the subgraph G_i has pathwidth less than k. Moreover, G_i has a path decomposition $\langle X_i^1, X_i^2, \ldots, X_i^{r_i}\rangle$ with $r_i \leq k$, such that $v_i^j \in X_i^j$ for $1 \leq j \leq r_i$, where $v_i^1, v_i^2, \ldots, v_i^{r_i}$ are the vertices of U_i indexed in cyclic order around the face F_i, clockwise or anti-clockwise.*

Lemma 13 ***(Grohe (2003))*** *The class of all graphs almost embeddable in a fixed surface S has linear local treewidth.*

Definition 14 *A graph G is h-almost-embeddable on a surface Σ if H is a h-apex graph of a graph that is almost embeddable on Σ.*

Theorem 15 ***(Robertson and Seymour (2003))*** *For any graph H, there is an integer $h \geq 0$ depending only on $|V(H)|$ such that any H-minor-free graph is a h-clique sum of a finite number of graphs that are h-almost-embeddable on some surfaces on which H cannot be embedded.*

Theorem 15 says that any H-minor-free graph G can be expressed as a "tree structure" of pieces, where each piece can be embedded on a surface on which H cannot be embedded after deleting at most h apex vertices.

Theorem 16 ***(Demaine et al. (2005))*** *For a fixed graph H, there is a constant c_H such that, for any integer $k \geq 1$ and for every H-minor-free graph G, the vertices of G (or the edges of G) can be partitioned into $k+1$ sets such that any k of the sets induce a graph of treewidth at most $c_H k$. Furthermore, such a partition can be found in polynomial time.*

Grohe et al. (2013) give a quadratic time algorithm that is faster for computing the clique-sum decomposition of H-minor-free graphs.

Definition 17 *Graph class $\mathcal{G}$ has truly sublinear treewidth with parameter λ where $0 < \lambda < 1$, if for every $\eta > 0$, there exists $\beta > 0$ such that for any graph $G \in \mathcal{G}$ and $X \subseteq V(G)$ the condition $tw(G \backslash X) \leq \eta$ yields that $tw(G) \leq \eta + \beta |X|^{\lambda}$.*

Lemma 18 ***(Fomin et al. (2011))*** *Let $\mathcal{G}_H$ be a class of graphs excluding a fixed graph H as a minor, then $\mathcal{G}_H$ has truly sublinear treewidth with $\lambda = \frac{1}{2}$.*

Theorem 19 *If $f_i \geq 0$ for all $i \in V$ where $|V| = n$, then for fixed k, there is an $O(q^{O(k)}kn)$-time algorithm for computing $\max_{\sigma \in [q]^V} \sum_{i \in V} f_i$ that achieves a solution of value at least $k/(k+2)$ optimal for H-minor-free graphs.*

Proof Our technique is similar to the proof in Demaine et al. (2005) of Theorem 16. Suppose the clique-sum decomposition of the input H-minor-free graph G is $G_1 \oplus G_2 \oplus \ldots \oplus G_t$ where each G_i ($1 \leq i \leq t$) is an h-almost embeddable graph. The join set J_i of the i-th clique-sum operation $(G_1 \oplus \ldots \oplus G_i) \oplus G_{i+1}$ is a subset of the apex set X_{i+1} of G_{i+1}. Our approximation algorithm takes the clique-sum decomposition as input, the apex set X_i of each G_i is given as part of the input clique-sum decomposition.

By the definition of the h-almost embeddable graphs, $G_i - X_i$ is almost embeddable on a bounded-genus surface where X_i contains at most h vertices. By lemma 13, $G_i - X_i$ has bounded local treewidth. From $i = 1$ to t, we choose a vertex $v_i \in G_i - X_i$ and construct a BFS tree T_i of $G_i - X_i$ rooted at v_i. Each vertex in $u \in G_i - X_i$ is labeled by the distance between u and v_i modulo $k+2$. After this step, we delete all the edges between levels labeled by $\ell \pmod{k+2}$ and the adjacent levels labeled by $\ell + 1 \pmod{k+2}$. Then the $G_i - X_i$ part is decomposed into several disjoint subgraphs with treewidth at most $c(k+2)$ for some constant $c > 0$. Since $J_{i-1} \subseteq X_i$, the vertices in J_{i-1} has already been labelled in $G_1 \oplus G_2 \oplus \ldots \oplus G_{i-1}$. We label the vertices in $X_i \backslash J_{i-1}$ arbitrarily by the integers from 0 to $k+1$. After the edge deletions, the obtained graphs G'_i are still H-minor-free for $1 \leq i \leq t$. By Lemma 18, the treewidth of G'_i is at most $ck + \beta |X_i|^{1/2} \leq ck + \beta h^{1/2} = O(k)$. It is known in Demaine et al. (2004) that $tw(G \oplus H) \leq \max\{tw(G), tw(H)\}$, thus we have $tw(G'_1 \oplus G'_2 \oplus \ldots \oplus G'_t) \leq \max\{tw(G'_1), \ldots, tw(G'_t)\}$. This shows that any H-minor-free graph can be transformed into a graph with treewidth bounded by $O(k)$ by deleting at most $|E|/k$ edges. Given a clique-sum decomposition, the vertex labelling and edge deletions can be finished in linear time.

Let A_i denote the vertices with labels other than $\ell \pmod{k+2}$ and $\ell + 1 \pmod{k+2}$ in G'_i and $B_i = V_i \backslash A_i$. Furthermore, we let $A = \bigcup_{i=1}^{t} A_i$ and $B = \bigcup_{i=1}^{t} B_i$. Since we only delete the edges between vertices by $\ell \pmod{k+2}$ and the adjacent levels labeled by $\ell + 1 \pmod{k+2}$, the functions f_i attached on vertices in A will not be undefined. The treewidth of the input graph is bounded by $O(k)$ after the edge deletions. We use the dynamic programming algorithm to compute the max-sum of functions f_i attached on vertices in A while ignoring the values of the functions f_i attached on vertices in B. The running time of the dynamic programming algorithm is $O(q^{O(k)}n)$.

Suppose the optimal solution is S_{OPT}, then there is at least one ℓ such that there is at most $\frac{2}{k+2}$ of the optimal solution falls into the vertices in B. If we repeat this computation for all $0 \leq \ell \leq k+1$, the optimal solution among these results is at least $\frac{k}{k+2}$ optimal. The time complexity of our algorithm is $O(q^{O(k)}kn)$. ■

2.4. Geometric graphs

We consider the d-dimensional Euclidean space $\mathbb{R}^d$. The distance between point i and point j is defined by the Euclidean distance, denoted by $\text{dist}(i,j) = ||i-j||_d$. Among various geometric objects, we only consider d-dimensional balls. We consider the following two scenarios.

Intersection Graph. Let d_i represent the diameter of the d-dimensional ball centered at point i for $i \in V$. An intersection graph $G(V,E)$ can be constructed by adding an edge between each pair of point i and point j where $\text{dist}(i,j) \leq (d_i + d_j)/2$.

Interference Graph. The model of interference graph is only defined for d-dimensional balls. The difference between interference graph and intersection graphs is that an interference graph is a directed graph while an intersection graph is undirected. There is an edge from the the center of ball i to the center of ball j if $\text{dist}(i,j) \leq d_i/2$.

The recognition of unit disk graphs is NP-hard Breu and Kirkpatrick (1998). For this reason, we assume that a geometric representation of the graph is always given as input. This assumption is reasonable in many applications. To defined the concepts of thickness and density, we first introduce the concepts of *slab decomposition* and *grid decomposition* in d-dimensional Euclidean space. A slab decomposition is partitioning the d-dimensional Euclidean space using parallel $(d-1)$-dimensional hyperplanes where the distances between neighboring hyperplanes are equal. A grid decomposition is recursively partitioning the d-dimensional Euclidean space using orthogonal slab decompositions where all the distances between neighboring hyperplanes are equal, until the space is partitioned into d-dimensional hypercubes of the same size. The distance between neighboring hyperplanes are set to $d_{\max}$ and $d_{\max}/2$ for intersection graphs and interference graphs respectively where $d_{\max}$ is the diameter of the maximum d-dimensional ball. By this decomposition, any vertex has only edges linked to vertices in the neighboring hypercubes.

Definition 20 *Given a slab decomposition, the thickness $th(G)$ of a geometric graph is determined by the maximum number of centers between any two neighboring hyperplanes. If a center falls precisely on a hyperplane $g(x_1,\ldots,x_d)=0$, it belongs to the subspace divided by the hyperplane where $g(x_1,\ldots,x_d)<0$.*

Definition 21 *Given a grid decomposition, the density $ds(G)$ of a geometric graph is determined by the maximum number of centers in any d-dimensional hypercube. If a center falls precisely on several hyperplanes $g_1,\ldots,g_m$, it belongs to the hypercube where $g_1<0,\ldots,g_m<0$.*

When the number of dimensions d is fixed, a slab decomposition or grid decomposition which gives the minimum thickness or density can be found in polynomial time.

Theorem 22 *Given the grid decomposition with density $ds(G)$, if $f_i \geq 0$ for all $i \in V$ where $|V|=n$, then for fixed k and d, there is an $O(q^{O((k+2)^{d-1}\cdot ds(G))}\cdot k^{d-1}\cdot n)$-time algorithm for computing $\max_{\sigma\in[q]^V}\sum_{i\in V} f_i$ that achieves a solution of value which is at least $(\frac{k}{k+2})^{d-1}$ optimal for d-dimensional geometric graphs with bounded density.*

For the first $d-1$ dimensions of the grid decomposition, we remove all the edges between the centers in hypercubes of levels congruent to $\ell_i \pmod{k+2}$ and the centers in hypercubes of levels

congruent to $\ell_i + 1 \pmod{k+2}$ for the i-th dimension and some integer $0 \leq \ell_i \leq k+1$. Then G is decomposed into several disjoint subgraphs each of which consists of $k+2$ levels of hypercubes in the first dimension, $k+2$ levels of hypercubes in the second dimension, ..., $k+2$ levels of hypercubes in the $(d-1)$-th dimension, all the levels of hypercubes in the d-th dimension. The pathwidths of these subgraphs are bounded by $O((k+2)^{d-1} \cdot ds(G))$. For each subgraph, we use the dynamic programming algorithm of Lemma 6 to maximize the sum of f_i attached on vertices in hypercubes that are in level 2 to level $k+1$ in the first dimension, level 2 to level $k+1$ in the second dimension, ..., level 2 to level $k+1$ in the $(d-1)$-th dimension and all levels in the d-th dimension while ignoring the values of f_i attached on all other vertices. By the pigeon hole principle, there is at least one tuple $(\ell_1, \ldots, \ell_{d-1})$ that the sum of the results on the subgraphs is at least $(\frac{k}{k+2})^{d-1}$ optimal. For computing the min-sum, we have a similar result as Theorem 9. More details of the proof are attached in Appendix D.

2.5. Graphs with bounded crossings per edge

Definition 23 *An embedding of a graph G on a surface S of genus g is a good embedding if all vertices of the graph are given as distinct points on S, no two edge crossings locate at the same point on S and for any edge, no vertex except the endpoints of the edge locate on the edge.*

Definition 24 *The crossing number φ of a graph G embedded on a surface S is the minimum over all good embeddings on S of the maximum over all edges e of the number of crossings of e.*

Several natural classes of graphs are graphs with bounded crossing parameter, such as geometric graphs with bounded density discussed above. It also includes graphs with bounded degree. Moreover, by the observation of Grigoriev and Bodlaender (2007), the class of graphs with bounded crossing parameter is not minor-closed. Therefore it generalizes the discussions about H-minor-free graphs and geometric graphs.

Theorem 25 *If $f_i \geq 0$ for all $i \in V$ where $|V| = n$, then for fixed k, there is an $O(q^{O(k)}kn)$-time algorithm for computing $\max_{\sigma \in [q]^V} \sum_{i \in V} f_i$ that achieves a solution of value at least $(k-\varphi-2)/k$ optimal for graphs with crossing parameter φ.*

To prove Theorem 25, we replace each crossing in the embedding of G by a new vertex. The obtained graph is denoted by G'. We construct a BFS tree T of G rooted at v. The level of a vertex is defined as the distance from the vertex to the root v on T. we remove every k-th level in T from G' together with its $\varphi - 1$ successive levels. The graph induced by the remaining vertices in T consists of several disjoint subgraphs with treewidth bounded by $O(k)$. The subgraph G_i has k_i levels of vertices in T where $k_i \in [k-\varphi, k]$. Then in the worst case, there is at least one partition of the graph that can achieve a solution at least $(k-\varphi-2)/k$ optimal. For computing the min-sum, we have a similar result as Theorem 9. More details of the proof are included in Appendix E.

Acknowledgments

This version was completed in February. In this version, we correct the error in the original version named *Beyond Baker's Technique*.

References

Gunnar Andersson, Lars Engebretsen, and Johan Håstad. A new way of using semidefinite programming with applications to linear equations mod p. *Journal of Algorithms*, 39(2):162–204, 2001.

Takao Asano. Approximation algorithms for max sat: Yannakakis vs. goemans-williamson. In *Theory of Computing and Systems, 1997., Proceedings of the Fifth Israeli Symposium on*, pages 24–37. IEEE, 1997.

Giorgio Ausiello. *Complexity and approximation: Combinatorial optimization problems and their approximability properties*. Springer, 1999.

Brenda S Baker. Approximation algorithms for np-complete problems on planar graphs. *Journal of the ACM (JACM)*, 41(1):153–180, 1994.

Hans L Bodlaender. A linear time algorithm for finding tree-decompositions of small treewidth. In *STOC*, pages 226–234. ACM, 1993.

Hans Leo Bodlaender. *Classes of graphs with bounded tree-width*. Department of Computer Science, University of Utrecht, 1986.

Yuri Boykov and Gareth Funka-Lea. Graph cuts and efficient nd image segmentation. *International journal of computer vision*, 70(2):109–131, 2006.

Yuri Boykov, Olga Veksler, and Ramin Zabih. Markov random fields with efficient approximations. In *IEEE computer society conference on Computer vision and pattern recognition*, pages 648–655. IEEE, 1998.

Yuri Boykov, Olga Veksler, and Ramin Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(11):1222–1239, 2001.

Heinz Breu and David G Kirkpatrick. Unit disk graph recognition is np-hard. *Computational Geometry*, 9(1):3–24, 1998.

David P Dailey. Uniqueness of colorability and colorability of planar 4-regular graphs are np-complete. *Discrete Mathematics*, 30(3):289–293, 1980.

Erik D Demaine and Mohammad Taghi Hajiaghayi. Diameter and treewidth in minor-closed graph families, revisited. *Algorithmica*, 40(3):211–215, 2004a.

Erik D Demaine and MohammadTaghi Hajiaghayi. Equivalence of local treewidth and linear local treewidth and its algorithmic applications. In *Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 840–849. SIAM, 2004b.

Erik D Demaine, MohammadTaghi Hajiaghayi, Naomi Nishimura, Prabhakar Ragde, and Dimitrios M Thilikos. Approximation algorithms for classes of graphs excluding single-crossing graphs as minors. *Journal of Computer and System Sciences*, 69(2):166–195, 2004.

Erik D Demaine, MohammadTaghi Hajiaghayi, and Ken-ichi Kawarabayashi. Algorithmic graph minor theory: Decomposition, approximation, and coloring. In *46th Annual IEEE Symposium on Foundations of Computer Science*, pages 637–646. IEEE, 2005.

Matt DeVos, Guoli Ding, Bogdan Oporowski, Daniel P Sanders, Bruce Reed, Paul Seymour, and Dirk Vertigan. Excluding any graph as a minor allows a low tree-width 2-coloring. *Journal of Combinatorial Theory, Series B*, 91(1):25–41, 2004.

Samuel Frederick Edwards and Phil W Anderson. Theory of spin glasses. *Journal of Physics F: Metal Physics*, 5(5):965, 1975.

David Eppstein. Subgraph isomorphism in planar graphs and related problems. In *SODA*, pages 632–640. Society for Industrial and Applied Mathematics, 1995.

David Eppstein. Diameter and treewidth in minor-closed graph families. *Algorithmica*, 27(3-4): 275–291, 2000.

Uriel Feige and Michel Goemans. Approximating the value of two power proof systems, with applications to max 2sat and max dicut. In *Third Israel Symposium on the Theory of Computing and Systems*, pages 182–189. IEEE, 1995.

Fedor V Fomin, Daniel Lokshtanov, Venkatesh Raman, and Saket Saurabh. Bidimensionality and eptas. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 748–759. SIAM, 2011.

Alan Frieze and Mark Jerrum. Improved approximation algorithms for max k-cut and max bisection. *Algorithmica*, 18(1):67–81, 1997.

Michel X Goemans and David Williamson. Approximation algorithms for max-3-cut and other problems via complex semidefinite programming. In *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 443–452. ACM, 2001.

Michel X Goemans and David P Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)*, 42 (6):1115–1145, 1995.

DM Greig, BT Porteous, and Allan H Seheult. Exact maximum a posteriori estimation for binary images. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 271–279, 1989.

Alexander Grigoriev and Hans L Bodlaender. Algorithms for graphs embeddable with few crossings per edge. *Algorithmica*, 49(1):1–11, 2007.

Martin Grohe. Local tree-width, excluded minors, and approximation algorithms. *Combinatorica*, 23(4):613–632, 2003.

Martin Grohe, Ken ichi Kawarabayashi, and Bruce A. Reed. A simple algorithm for the graph minor decomposition - logic meets structural graph theory. In *SODA*, pages 414–431. SIAM, 2013.

F Hadlock. Finding a maximum cut of a planar graph in polynomial time. *SIAM Journal on Computing*, 4(3):221–225, 1975.

Johan Håstad. Some optimal inapproximability results. *Journal of the ACM (JACM)*, 48(4):798–859, 2001.

Dorit S Hochbaum and Wolfgang Maass. Approximation schemes for covering and packing problems in image processing and vlsi. *Journal of the ACM (JACM)*, 32(1):130–136, 1985.

John Hopcroft and Robert Tarjan. Efficient planarity testing. *Journal of the ACM (JACM)*, 21(4): 549–568, 1974.

Harry B Hunt III, Madhav V Marathe, Venkatesh Radhakrishnan, Shankar S Ravi, Daniel J Rosenkrantz, and Richard E Stearns. Nc-approximation schemes for np-and pspace-hard problems for geometric graphs. *Journal of Algorithms*, 26(2):238–274, 1998.

Howard Karloff and Uri Zwick. A 7/8-approximation algorithm for max 3sat? In *Proceedings of 38th Annual Symposium on Foundations of Computer Science*, pages 406–415. IEEE, 1997.

Sanjeev Khanna and Rajeev Motwani. Towards a syntactic characterization of ptas. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 329–337. ACM, 1996.

Subhash Khot and Nisheeth K Vishnoi. On the unique games conjecture. In *FOCS*, volume 5, page 3, 2005.

Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O'Donnell. Optimal inapproximability results for max-cut and other 2-variable csps? *SIAM Journal on Computing*, 37(1):319–357, 2007.

Philip N Klein. A linear-time approximation scheme for tsp in undirected planar graphs with edge-weights. *SIAM Journal on Computing*, 37(6):1926–1952, 2008.

Vladimir Kolmogorov and Ramin Zabih. Multi-camera scene reconstruction via graph cuts. In *ECCV 2002*, pages 82–96. Springer, 2002.

Frank R Kschischang, Brendan J Frey, and H-A Loeliger. Factor graphs and the sum-product algorithm. *IEEE Transactions on Information Theory*, 47(2):498–519, 2001.

Shiro Matuura and Tomomi Matsui. 0.863 approximation algorithm for max dicut. *Approximation, Randomization, and Combinatorial Optimization: Algorithms and Techniques Lecture Notes in Computer Science*, 2129:138–146, 2001.

Christos H Papadimitriou and Mihalis Yannakakis. Optimization, approximation, and complexity classes. *Journal of computer and system sciences*, 43(3):425–440, 1991.

Judea Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann, 1988.

Renfrey Burnard Potts. Some generalized order-disorder transformations. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 48, pages 106–109. Cambridge Univ Press, 1952.

Neil Robertson and Paul D Seymour. Graph minors. iii. planar tree-width. *Journal of Combinatorial Theory, Series B*, 36(1):49–64, 1984.

Neil Robertson and Paul D Seymour. Graph minors. xiii. the disjoint paths problem. *Journal of combinatorial theory, Series B*, 63(1):65–110, 1995.

Neil Robertson and Paul D Seymour. Graph minors. xvi. excluding a non-planar graph. *Journal of Combinatorial Theory, Series B*, 89(1):43–76, 2003.

Neil Robertson and Paul D Seymour. Graph minors. xx. wagner's conjecture. *Journal of Combinatorial Theory, Series B*, 92(2):325–357, 2004.

Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. Grabcut: Interactive foreground extraction using iterated graph cuts. In *ACM Transactions on Graphics (TOG)*, volume 23, pages 309–314. ACM, 2004.

Jian Sun, Nan-Ning Zheng, and Heung-Yeung Shum. Stereo matching using belief propagation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(7):787–800, 2003.

Luca Trevisan, Gregory B Sorkin, Madhu Sudan, and David P Williamson. Gadgets, approximation, and linear programming. *SIAM Journal on Computing*, 29(6):2074–2097, 2000.

Erik Jan van Leeuwen. Approximation algorithms for unit disk graphs. In *Graph-theoretic concepts in computer science*, pages 351–361. Springer, 2005.

Klaus Wagner. Über eine eigenschaft der ebenen komplexe. *Mathematische Annalen*, 114(1):570–590, 1937.

Jonathan S Yedidia, William T Freeman, Yair Weiss, et al. Generalized belief propagation. In *NIPS*, volume 13, pages 689–695, 2000.

Appendix A. Related Work

Our method strongly generalizes the technique in Baker (1994) which is a method for designing PTASs for NP-hard optimization problems on planar graphs based on layerwise decomposition. It is also independently discovered by Hochbaum and Maass (1985), called *Shifting Strategy*. It gives PTASs to maximum independent set, minimum vertex cover, minimum dominating set and some other NP-hard optimization problems. In Eppstein (1995), it is used to give a PTAS for subgraph isomorphism, connectivity and shortest path problems on planar graphs and bounded genus graphs. Khanna and Motwani (1996) proves that when the corresponding incidence graphs of MP-SAT, TMAX and TMIN formulas are planar, then there are PTASs for computing their MAX-SAT. Eppstein (2000) generalizes the PTASs from graphs with bounded outerplanarity to graphs with bounded local treewidth. Eppstein (2000) further characterizes that a minor-closed class of graphs has bounded local treewidth if and only if it does not contain all apex graphs. The bound of local treewidth for apex-minor-free graphs in Eppstein (2000) is $2^{2^{O(r)}}$, and $O(r)$ for bounded-genus graphs which are also apex-minor-free graphs. This bound is improved in Demaine and Hajiaghayi (2004a) to $2^{O(r)}$ and further improved by Demaine and Hajiaghayi (2004b) to the optimal linear

bound $O(r)$. The linear bound of local treewidth for apex-minor-free graphs gives PTASs to several optimization problems on apex-minor-free graphs with time complexity $2^{O(1/\epsilon)}n^{O(1)}$. In Klein (2008), Baker's technique is applied to the travelling salesman problems (TSP) on undirected planar graphs and achieves a linear time approximation scheme. The technique of Baker (1994) has also been applied to obtain PTASs for optimization problems on H-minor-free graphs, which relies heavily on a deep structural theorem Robertson and Seymour (2003) of Robertson and Seymour's graph minor theory. Grohe (2003) proves that there is a PTAS for computing minimum vertex cover, minimum dominating set and maximum independent set on H-minor-free graphs. However, this algorithm is existential since constructing the tree decomposition relies on the polynomial time membership test of minor-closed properties, which is non-constructive Robertson and Seymour (1995). Demaine et al. (2005) provide an algorithm to construct a decomposition with additional properties proved in DeVos et al. (2004) in $O(n^c)$ time where c is a constant depending on the size of the excluded minor H. This makes Robertson and Seymour's result constructive. Grohe et al. (2013) give a quadratic time algorithm for computing graph minor decompositions.

MAX-SAT is APX-hard so that there is no PTAS in general unless P = NP. Moreover, although the decision problem 2-SAT can be decided in polynomial time, MAX-2SAT is APX-hard Ausiello (1999). Feige and Goemans (1995) provide a 0.931-approximation algorithm for MAX-2SAT, improving the approximation ratio of 0.87856 in Goemans and Williamson (1995). Trevisan et al. (2000) give a 0.801-approximation algorithm for MAX-3SAT. If the formula is satisfiable, then there is a 7/8-approximation algorithm for MAX-3SAT Karloff and Zwick (1997). Asano (1997) gives a 0.770-approximation algorithm for MAX-SAT. Håstad (2001) shows that MAX-3SAT is hard to approximate within 7/8 and MAX-2SAT is hard to approximate within $73/74 \approx 0.986$.

On planar graphs, MAX-CUT is dual to the route inspection problem so that it is polynomial-time solvable Hadlock (1975). In general, this problem is APX-hard Papadimitriou and Yannakakis (1991). The best known approximation ratio achieved by polynomial time algorithm for computing MAX-CUT is $\alpha = \frac{2}{\pi}\min_{0\le\theta\le\pi}\frac{\theta}{1-\cos\theta} \approx 0.878$ discovered by Goemans and Williamson (1995). Khot et al. (2007) show that this is the best possible approximation ratio for MAX-CUT if the unique game conjecture Khot and Vishnoi (2005) is true. Without such unproved assumptions, Trevisan et al. (2000) show that it is NP-hard to approximate MAX-CUT within $\frac{16}{17} \approx 0.941$ optimal. Goemans and Williamson (1995) obtain a 0.796-approximation algorithm for MAX-DICUT. Feige and Goemans (1995) improve this result to 0.859. Matuura and Matsui (2001) give a 0.863-approximation algorithm for the MAX-DICUT problem. Frieze and Jerrum (1997) prove that the randomized approximation ratio α_k of MAX k-CUT satisfies $\alpha_k \approx \frac{k-1}{k} + 2k^2 \ln k$ when $k \to \infty$. This implies $\alpha_2 \ge 0.878567$, $\alpha_3 \ge 0.800217$, $\alpha_4 \ge 0.850304$ and so on. Goemans and Williamson (2001) use complex semidefinite programming and an extension of the random hyperplane technique to obtain an approximation algorithm for MAX-3-DICUT with approximation ratio 0.79373 and MAX 3-CUT with approximation ratio $\frac{7}{12} + \frac{3}{4\pi^2}\arccos^2(-1/4) \approx 0.83601$. Andersson et al. (2001) show that it is NP-hard to approximate Max 2-Lin-Mod-3 with $\alpha > 17/18 \approx 0.944$ unless P = NP, which implies that MAX 3-CUT and MAX-DICUT also cannot be approximated within $17/18$ unless P = NP.

Appendix B. Definitions and Notations

In this section we introduce some notations and concepts that will be used throughout this paper. Let G be a graph whose vertex set is $V(G)$ and edge set is $E(G)$. A graph G' is a *subgraph*

of G if $V(G') \subseteq V(G)$ and $E(G') \subseteq E(G)$. The subgraph G' is called a *induced subgraph* of $V(G')$ if $E(G') = \{(u,v) \in E(G) \mid u,v \in V(G')\}$. The subgraph induced by V' is denoted by $G[V']$. For $v \in V$, the neighborhood of v is denoted by $N(v)$ where $N(v) = \{u \mid (u,v) \in E\}$. For $u,v \in V(G)$, the distance $\text{dist}_G(u,v)$ is the length of shortest path between u and v. The r-neighborhood of vertex $v \in V$ is defined as $N_r(v) = \{u : \text{dist}_G(u,v) \leq r\}$.

Approximation schemes. A PTAS is an algorithm $\mathcal{A}(I,\epsilon)$ which takes an instance I of an optimization problem and a parameter $\epsilon > 0$ and runs in time $n^{O(f(1/\epsilon))}$ which produces a solution that is at least $(1-\epsilon)$ optimal for maximization and at most $(1+\epsilon)$ optimal for minimization. A PTAS with running time $f(1/\epsilon) \cdot n^{O(1)}$ is called an efficient polynomial time approximation scheme (EPTAS). An EPTAS where $f(1/\epsilon)$ is polynomial in $1/\epsilon$ is called a fully polynomial time approximation scheme (FPTAS).

Planar graph. A graph G is a planar graph if it can be embedded into the two-dimensional plane that no pair of edges will cross with each other. Given a planar graph, its planar embedding G can be generated in linear time Hopcroft and Tarjan (1974). A planar graph is an outerplanar graph if it has a planar embedding where all the nodes are on the exterior face. Given a planar embedding of a planar graph, a node is at level 1 if it is on the exterior face. When all the level-1 nodes are deleted from the planar embedding, the nodes on the exterior face are called level-2 nodes. By this induction the level-k nodes can be defined. A planar graph is a k-outerplanar graph if it has a planar embedding with no nodes of level more than k.

Graph minor. Given a graph G, if graph H can be reduced from a subgraph of G by a sequence of edge contractions, then H is a *minor* of G, denoted by $H \leq_m G$. We can see that $H \leq_m G$ if and only if there is a mapping $h : V_H \to 2^{V_G}$ such that $G[h(x)]$ is a connected subgraph of G for all $x \in V_H$, $h(x) \cap h(y) = \emptyset$ for all $x \neq y \in V_H$ and, for every $(x,y) \in E_H$ there exists an edge $(u,v) \in E_G$ such that $u \in h(x)$ and $v \in h(y)$. The mapping h is called a *witness* of $H \leq_m G$.

A class $\mathcal{C}$ of graphs is minor-closed if and only if for all $G \in \mathcal{C}$ and $H \leq_m G$ we have $H \in \mathcal{C}$. We say $\mathcal{C}$ is nontrivial if $\mathcal{C}$ does not contain all the graphs. A class $\mathcal{C}$ of graphs is H-minor-free if $H \not\leq_m G$ for all $G \in \mathcal{C}$. Then we call H an excluded minor of $\mathcal{C}$. Robertson and Seymour's Graph Minor Theorem Robertson and Seymour (2004), which proves Wagner's conjecture, demonstrates that the undirected graphs partially ordered by the graph minor relationship form a well-quasi-ordering. This implies that every minor-closed class of graphs can be characterized by a finite set of forbidden minors.

Tree decomposition and path decomposition. The concept of tree decomposition is introduced by Robertson and Seymour (1984), on which the treewidth can be defined to measure the similarity between a graph and a tree, has turned out to be significant for structural graph theory and graph algorithms.

Definition 26 *A tree decomposition of an undirected graph $G(V,E)$ is a pair*

$$(\{X_i | i \in I\}, T = (I,F))$$

with $\{X_i | i \in I\}$ a family of subsets of V, one for each node of T. T is a tree such that

1. $\bigcup_{i \in I} X_i = V$,

2. *for all edges* $\{v, w\} \in E$*, there exists an* $i \in I$ *with* $v \in X_i$ *and* $w \in X_i$*,*
3. *for all* $i, j, k \in I$*: if* j *is on the path from* i *to* k *in* T*, then* $X_i \cap X_k \subseteq X_j$*.*

Each node of the tree decomposition T is called a *bag*. The third property of tree decomposition guarantees that for every $v \in V$, $\{X_i : v \in X_i, i \in I\}$ induces a connected subtree of T.

Definition 27 *The treewidth of a tree decomposition* $(\{X_i | i \in T\}, T = (I, F))$ *is* $\max_{i \in I} |X_i| - 1$*. The treewidth of a graph* G*, denoted by* $tw(G)$*, is the minimum treewidth over all tree decompositions of* G*.*

A tree decomposition of width equal to the treewidth is called an optimal tree decomposition. Computing the treewidth for graph G is NP-complete. But given a graph G, deciding whether the treewidth of G is at most a fixed constant k can be decided in linear time by Bodlaender's algorithm Bodlaender (1993). If the answer is yes, then an optimal tree decomposition of G can be constructed in linear time (but exponential in k). The following lemma includes some well-known facts about treewidth.

Lemma 28 *Let* $(\{X_i | i \in I\}, T = (I, F))$ *be a tree decomposition of graph* G*. Then the following holds:*

1. *If* $X \subseteq V(G)$ *is a clique, then there is an* $i \in I$ *that* $X \subseteq X_i$*.*
2. *Let* G, H *be graphs such that* $V(G) \cap V(H)$ *is a clique in both* G *and* H*. Then it holds that* $tw(G \cup H) = \max\{tw(G), tw(H)\}$*.*
3. *For any* $X \subseteq V(G)$*. Then* $tw(G) \leq tw(G \backslash X) + |X|$*.*
4. *Let* G, H *be graphs such that* $H \leq_m G$*. Then* $tw(H) \leq tw(G)$*.*

Definition 29 *A tree decomposition is called a path decomposition if* $T = (I, F)$ *is a path. The pathwidth of a graph* G*, denoted by* $pw(G)$*, is the minimum width over all path decompositions of* G*.*

Local treewidth. The concept of local treewidth is first introduced by Eppstein (2000) as a generalization of treewidth. The local treewidth of graph G is a function that maps an integer $r \in \mathbb{N}$ to the maximum treewidth of the subgraph of G induced by the r-neighborhood of any vertex in G, formally defined as follows.

Definition 30 *The local treewidth of graph* $G(V, E)$ *is a function defined as*

$$ltw^G(r) = \max\left\{tw(G[N_r(i)]) : i \in V\right\}$$

where $G[N_r(i)]$ *is the subgraph of* G *induced by* $N_r(i)$*.*

Definition 31 *A class* $\mathcal{C}$ *of graphs has bounded local treewidth if there is a function* $f : \mathbb{N} \to \mathbb{N}$ *such that* $ltw^G(r) \leq f(r)$ *for all* $G \in \mathcal{C}$*,* $r \in \mathbb{N}$*.* $\mathcal{C}$ *has linear local treewidth if there is a* $\lambda \in \mathbb{R}$ *such that* $ltw^G(r) \leq \lambda \cdot r$ *for all* $G \in \mathcal{C}$*,* $r \in \mathbb{N}$*.*

It is well known that the treewidth of a $k \times k$ grid is k, so planar graphs do not have bounded treewidth. But for each $i \in V$ of any planar graph G, the subgraph induced by its r-neighborhood $N_r(v)$ has bounded treewidth. Moreover, the k-outerplanar graph has treewidth at most $3k - 1$ Bodlaender (1986).

Graphs on surfaces. A surface Σ is a compact and connected 2-manifold without boundary. An embedding refers to a 2-cell embedding. An embedding of a graph G on a surface Σ is a representation of G on surface Σ that maps vertices and edges of G to points and arcs on Σ while preserving their correlations in G. Every face bounded by arcs is homomorphic to an open disk, that is, $\{(x, y)|x^2 + y^2 < 1\}$. Therefore, an embedding of graph on a surface is drawing the graph on the surface such that the edges only intersect at their endpoints. A graph has bounded genus if it can be embedded on a bounded-genus surface. The genus of a graph is the minimum genus of all surfaces on which it can be embedded.

Clique sum. The clique sum operation is a significant technical tool of graph minor theory, which originates from characterizations of $K_{3,3}$-minor-free and K_5-minor-free graphs by Wagner (1937). The clique-sum operation is a way of combining two graphs by identifying their cliques. Suppose G_1 and G_2 are two graphs, $W_1 \subseteq V(G_1)$ and $W_2 \subseteq V(G_2)$ are two cliques of the same size. The clique sum of G_1 and G_2, denoted by $G_1 \oplus G_2$, is a graph by identifying W_1 and W_2 through a bijection $h : W_1 \to W_2$, and then possibly deleting some of the clique edges. The subgraph induced by the clique vertices in $G_1 \oplus G_2$ is called the *join set*. The clique is called a k-sum if $|W_1| = |W_2| = k$, denoted by $G_1 \oplus_k G_2$. Since there are many possible bijections between vertices of W_1 and W_2, there are also many possible results for $G_1 \oplus G_2$.

The clique-sum operation plays an important role in the core of Robertson and Seymour's graph minor theory. The deepest structural theorem Robertson and Seymour (2003) of graph minor theory states that any H-minor-free graph can be decomposed into a collection of graphs each of which can be embedded into a bounded-genus surface by deleting a bounded number of apex vertices where the number only depends on $V(H)$. These so-called h-almost embedded graphs are combined in a tree structure by clique-sum operations. The clique-sum decomposition Demaine et al. (2005); Grohe et al. (2013) is the building block by which the approximation algorithms for H-minor-free graphs can be achieved.

Appendix C. Graphs with bounded local treewidth

The following theorem gives a precise characterization of graphs with bounded local treewidth.

Theorem 32 ***(Eppstein (2000))*** *Let $\mathcal{F}$ be a minor closed family of graphs. Then $\mathcal{F}$ has bounded local treewidth if and only if $\mathcal{F}$ does not contain all apex graphs.*

A graph is an apex graph if it has a vertex whose removal results in a planar graph. For example, K_5 and $K_{3,3}$ are apex graphs. Theorem 32 shows that a graph G has bounded local treewidth if and only if it is apex-minor-free.

Theorem 33 ***(Demaine and Hajiaghayi (2004b))*** *Any apex-minor-free graph has linear local treewidth.*

The linearity of local treewidth has been proved for some subclasses of apex-minor-free graphs before Demaine and Hajiaghayi (2004b), such as bounded-genus graphs Eppstein (2000) and single-crossing-minor-free graphs Demaine et al. (2004).

Proof [Proof of Theorem 11] Choosing any vertex $v \in V$ as root, construct a BFS tree T rooted at v. The layer of vertices is defined as its distance to v. Moreover, the set of vertices from layer i to layer j is denoted by

$$L_v^G[i,j] = \{u \in V | i \leq \text{dist}(u,v) \leq j\}.$$

If $i > j$, $L_v^G[i,j] = \emptyset$. For any $i \leq j$, $L_v^G[i,j]$ has bounded local treewidth. This is because if we obtain a minor H of G by contracting the subgraph of G induced by $L_v^G[0,i-1]$ to a single vertex v', $L_v^G[i,j] \subseteq L_{v'}^H[1,j-i+1]$. Since G is apex-minor-free, H is also apex-minor-free. Therefore, H has bounded local treewidth. Then we have $tw(L_v^G[i,j]) = O(j-i+1)$, which implies any subgraph induced by consecutive k levels of vertices in T has treewidth bounded by $O(k)$.

We delete all the edges between levels congruent to $\ell \pmod{k+2}$ and $\ell+1 \pmod{k+2}$ for some integer ℓ that $0 \leq \ell \leq k+1$. Then G is decomposed into several disjoint subgraphs $G_1, \ldots, G_t$. For each G_i where $1 \leq i \leq t$, we construct a tree decomposition T_i. We still use the dynamic programming algorithm on T_i to maximize the sum of f_i attached on non-boundary vertices while ignoring the values of f_i on the boundary of G_i. Since $tw(G_i) = O(k)$, the time complexity of the dynamic programming is $O(q^{O(k)}n)$. By a similar argument of Theorem 7, there is at least one ℓ that the combination of the solutions on the disjoint subgraphs is at least $\frac{k}{k+2}$ optimal of the optimal solution. ■

Appendix D. Geometric graphs

Among various geometric objects, we only consider d-dimensional balls. Our results also hold for other geometric objects no matter it is convex or not. To do this, we choose the midpoint of the line determined by two points of the longest distance on the object as the center and the longest distance from the center to any point of the object as the radius. Then the terminologies of d-dimensional balls can also be used for these objects. Their geometric graphs can be constructed in a similar way.

A geometric graph has several important parameters related to the performances of our approximation algorithms, listed as follows. The concept of λ-precision is presented in Hunt III et al. (1998). The concepts of *thickness* and *density* are presented in van Leeuwen (2005). The concept of λ-precision is defined as follows.

Definition 34 *For $\lambda > 0$, a geometric graph is a λ-precision geometric graph if the centers of any two d-dimensional balls are at least λ apart.*

A λ-precision geometric graph has density $\Theta(\frac{1}{\lambda^d})$ in d-dimensional Euclidean space.

Proof [Proof of Theorem 22] For the first dimension of the grid decomposition, we remove all the edges between the centers in hypercubes of levels congruent to $\ell_1 \pmod{k+2}$ and the centers in hypercubes of levels congruent to $\ell_1+1 \pmod{k+2}$ for some $0 \leq \ell_1 \leq k+1$. Then G is decomposed into several disjoint subgraphs with $k+2$ levels of hypercubes in the first dimension, all the levels of hypercubes in the second dimension, ..., all the levels of hypercubes in the d-th dimension, denoted by $G_1^1, \ldots, G_{t_1}^1$.

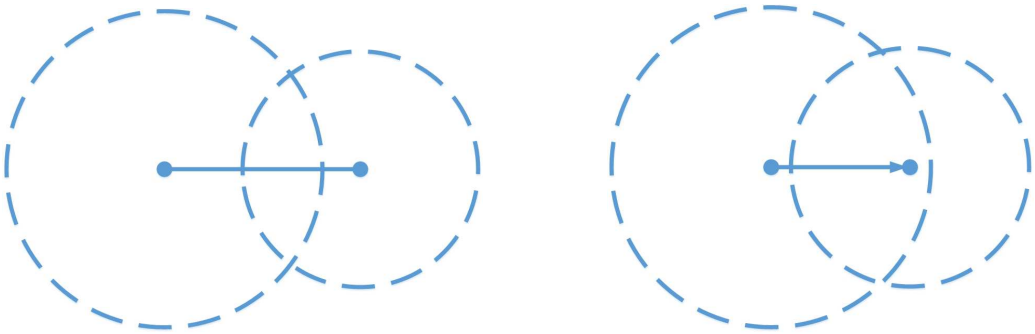

(*a*) Intersection graph (*b*) Interference graph

Figure 3: Geometric graphs

For the second dimension of G_i^1 for some $1 \leq i \leq t_1$, we remove all the edges between the centers in hypercubes of levels congruent to $\ell_2 \pmod{k+2}$ and the centers in hypercubes of levels congruent to $\ell_2 + 1 \pmod{k+2}$ for some $0 \leq \ell_2 \leq k+1$. Then G_i^1 is decomposed into several disjoint subgraphs with $k+2$ levels of hypercubes in the first dimension, $k+2$ levels of hypercubes in the second dimension, all the levels of hypercubes in the third dimension, ..., all the levels of hypercubes in the d-th dimension, denoted by $G_1^2, \ldots, G_{t_2}^2$.

We repeat this procedure until for the $(d-1)$-th dimension of G_i^{d-2} for some $1 \leq i \leq t_{d-2}$, we remove all the edges between the centers in hypercubes of levels congruent to $\ell_{d-1} \pmod{k+2}$ and the centers in hypercubes of levels congruent to $\ell_{d-1}+1 \pmod{k+2}$ for some $0 \leq \ell_{d-1} \leq k+1$. Then G_i^{d-2} is decomposed into several disjoint subgraphs each of which consists of $k+2$ levels of hypercubes in the first dimension, $k+2$ levels of hypercubes in the second dimension, ..., $k+2$ levels of hypercubes in the $(d-1)$-th dimension, all the levels of hypercubes in the d-th dimension, denoted by $G_1^{d-1}, \ldots, G_{t_{d-1}}^{d-1}$. For each subgraph G_i^{d-1}, it satisfies that

$$th(G_i^{d-1}) \leq (k+2)^{d-1} \cdot ds(G).$$

for all $1 \leq i \leq t_{d-1}$. Moreover, the pathwidth of G_i^{d-1} is at most $2 \cdot th(G_i^{d-1}) - 1$ where each bag of the path decomposition contains two neighboring hypercubes.

For each G_i^{d-1}, we use the dynamic programming algorithm of Lemma 6 to maximize the sum of functions f_i attached on vertices in hypercubes that are in level 2 to level $k+1$ in the first dimension, level 2 to level $k+1$ in the second dimension, ..., level 2 to level $k+1$ in the $(d-1)$-th dimension and all levels in the d-th dimension while ignoring the values of functions f_i attached on all other vertices. The time complexity of the dynamic programming on each G_i^{d-1} is $O\big(q^{O\left((k+2)^{d-1} \cdot ds(G)\right)} n\big)$.

By the pigeon hole principle, there is at least one tuple $(\ell_1, \ldots, \ell_{d-1})$ that the sum of the results of dynamic programming on each G_i^{d-1} is at least $(\frac{k}{k+2})^{d-1}$ of the optimal solution. For each possible $(\ell_1, \ldots, \ell_{d-1})$, we need to compute the dynamic programming for all the G_i^{d-1}. Therefore, the time complexity of our algorithm is $O\big(q^{O\left((k+2)^{d-1} \cdot ds(G)\right)} \cdot k^{d-1} \cdot n\big)$. ■

This result gives us a PTAS for computing the max-sum of functions f_i on geometric graphs. Given an error $\epsilon > 0$, we choose k as

$$k = \left\lceil \frac{2(1-\epsilon)^{\frac{1}{d-1}}}{1-(1-\epsilon)^{\frac{1}{d-1}}} \right\rceil,$$

then the solution achieved by our algorithm is at least $(1-\epsilon)$ optimal.

Appendix E. Graphs with bounded crossings per edge

Proof [Proof of Theorem 25] We obtain a planar graph $G' = (V', E')$ by replacing each edge crossing of G by a new vertex. Construct a breadth first search tree T of G, rooted at $v \in V$. The level of a vertex is defined as the distance from the vertex to the root v on T. Suppose $\varphi < k$. We remove every k-th level in T from G' together with its $\varphi - 1$ successive levels. Then G' is decomposed into several subgraphs $\mathcal{H} = \{H_1, \ldots, H_t\}$, where each $H_i = (N_i, E_i)$ that $1 \leq i \leq t$ is induced by $k - \varphi$ consecutive levels of vertices of in T of G'.

Let $V_i = N_i \cap V$ and $G_i = G[V_i]$, which is the subgraph of G induced by V_i. Since the number of crossings per edge is at most φ and φ consecutive levels of vertices are removed from G' for every k levels, thus after the removal all the subgraphs G_i are disjoint. For each G_i which consists of $k_i \in [k-\varphi, k]$ levels of vertices in T, denoted by $V_{i,1}, \ldots, V_{i,k_i}$. Note that only the vertices in $V_{i,1}$ are possible to be adjacent to vertices in $V_{i-1,k_{i-1}}$ and only the vertices in V_{i,k_i} are possible to be adjacent to the vertices $V_{i+1,1}$ in G. It is proved in Grigoriev and Bodlaender (2007) that each G_i has treewidth $O(k)$. Let $A_i = \bigcup_{j=2}^{k_i-1} V_{i,j}$ and $B_i = V_i - A_i$. For each G_i we construct a tree decomposition T_i of treewidth $O(k)$. And we use a dynamic programming on T_i to maximize $\sum_{j \in A_i} f_j$ while ignoring the values of f_i defined on the vertices in B_i. After the computation, we combine the solutions on each G_i to form the final solution. By similar arguments, there is at least one partition can achieve a solution which is at least $(k-\varphi-2)/k$ optimal. ■